Mathematical Epistemologies at Work [1]

RICHARD NOSS

In mathematical terms, there is a celebrated tension between forms of discourse and cognition that are delicately tuned to cultural practices and those that are focused explicitly on mathematics *per se*, recognisable by its symbolic forms and epistemological structures This tension parallels (and is perhaps derived from) the epistemological duality of mathematical thought as both tool and object, simultaneously a component of pragmatic activity and theoretical endeavour

The preparation of this article has afforded an opportunity to reflect retrospectively on this duality and on a corpus of research in which I and my colleagues have been involved, spanning a variety of sub-fields and a couple of decades. I hope it is not too fanciful to impose upon this work a narrative that was not necessarily evident to any of us while we were engaged upon it. Here is a first outline of that narrative

I begin with a pervasive finding that arises from investigations with (mainly young) people expressing mathematical ideas with computers. These studies led to a series of thoughts concerning the generation of mathematical meanings that nagged away until the early nineteen-nineties, when Celia Hoyles and myself began to formulate a theoretical framework for describing the phenomena we encountered Shortly after this, we had the opportunity to work in a variety of settings with the broad common aim of elaborating the mathematics used in working practices.

I shall then illustrate how these studies began to throw light on some fundamental questions, particularly concerning the nature of mathematical practices, and encouraged us to investigate further the problem of mathematical meaning from both cognitive and socio-cultural perspectives. This effort has led to some general principles about the design of mathematical activity systems for learning and, in particular, the rather special role that digital technologies may play within them. Thus, perhaps fittingly but probably over-ambitiously, I will conclude where I began, with the assertion that digital technologies can play an unusually powerful role in helping to understand and reshape the nature of mathematical sense-making.

I would like to make two general observations at the outset. The first concerns my wish to consider both cognitive and social dimensions. To steer a course between these two approaches is not easy, not least because proponents of each often ignore the work of the other, or denounce as mere eclecticism any attempt at synthesis (there are important exceptions to this: see, for example, Cobb and Bowers, 1999; Kieran, Forman and Sfard, 2001) One organising idea for thinking about this apparent dichotomy has been suggested to me by Andy diSessa who distinguishes between phenomena that are *distally* and *proximally* social

Much of what I have to say comes from a recognition that many phenomena concerned with mathematical meaning are proximally social, in that they manifestly involve social and cultural relations between people and within communities. But I also recognise that many facets of human thought are only distally social; while it is true that what I think, and the techniques I use for thinking and communicating are shaped both socially and culturally, I think in ways that are structured by my personal cognitive history at least as strongly as by the socio-cultural relationships in which I find myself embedded.

No attempt to understand how mathematics is learned by human beings can afford to ignore this essentially cognitive element, any more than it can afford to ignore the social and cultural relations in which cognitive activity is embedded. Thus, in what follows, I hope to illustrate not only that such a perspective need not necessarily lapse into eclecticism, but rather that co-ordination of the two approaches provides a possible and even necessary methodological stance.

The second observation concerns the title of this article. I recognise that it is bad form to tell a joke and then explain it Forgive me then, if I explain the *double entendre* in the title I want to talk about mathematical epistemology as it is found *in work*, to understand how mathematics is used and how it is conceived by participants in their cultural practices. But I also want to talk about mathematical epistemology as a crucial element *at work* in learning situations; how mathematics education researchers can develop not just new approaches to teaching, but new mathematical epistemologies that are more learnable and, at least for all but the few, more expressive

Insights from observations of activities with computers

Over some two decades, Celia Hoyles and I have engaged in studies of children and adults interacting with computational systems designed to afford mathematical expression. Throughout this time, we have noticed an interesting phenomenon, which we can simplistically characterise as follows: learners are often able to express themselves in terms that might be considered abstract, yet which seem to be bound tightly into the tools and symbols of the computational world. Learners can, in other words, say and do things with suitably-designed systems that they may be unable to say or do without them - and they can often do so in ways that are interestingly different from conventional means. Let me give an example of this phenomenon from the recently-completed doctoral thesis of my colleague, Lulu Healy. Healy (2002) reports the results of an investigation into children's understanding of reflection and symmetry, in which she designed, constructed and evaluated two learning systems, one based

For the Learning of Mathematics **22**, 2 (July, 2002)
FLM Publishing Association, Kingston, Ontario, Canada

on Cabri and the other on a Logo-based toolset.

Within this toolset was one particular tool, named *meet*, that simulated the action of turtles moving successively closer (their 'speeds' adjusted accordingly) in order to construct a new turtle at the point at which their paths intersect Figure 1 shows three snapshots of the general *meet* tool in action.

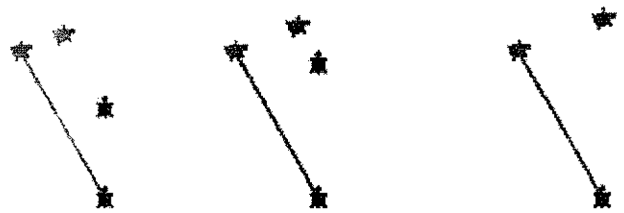

Figure 1: The 'meet' tool in action: copies of each of the turtles at the ends of the line segment inch forward until they meet and become one

In the special case in which the two turtles are initially placed in a (reflective) symmetric configuration, they will, of course, meet on the line of symmetry: surprisingly or not, Healy's twelve-year-old students all appeared to acknowledge this intuitively The *meet* tool, therefore, was designed to tap into this intuition and afford students a way both to construct symmetrical figures (see Figure 2), as well as to justify hypotheses that one turtle was the reflection of another in a given line of symmetry.

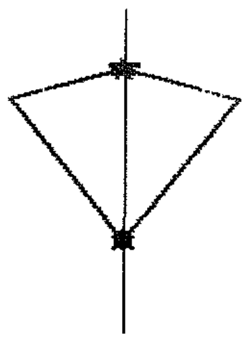

Figure 2: Using the 'meet' tool to complete a symmetrical figure

I will suppress all of the interesting detail in order to focus on the ways that children expressed their ideas about symmetry during interactions with this tool. As an example, Aimee said:

> Every turtle has its own reflection turtle [with] the same distance away from the mirror and the same angle, except for lefts and rights.

Notice that 'except for lefts and rights' states precisely what it means *within* the Logo setting - under reasonable conditions, a figure is symmetrical if and only if it consists of two parts which are the same with *left* and *right* swapped. Aimee's expression captures the properties of the figures drawn on the screen; but it also focuses on the relationships that associate an object with its image (and line of symmetry), suggesting a view of reflection as a transformation of a set of turtles (points)

Statements such as Aimee's could be used to reason about reflection, i.e to generate implied knowledge. Here, for example, is Jodie's reasoning:

> It's wrong. If it was symmetrical, the turtles would meet on the mirror This one looks the same distance as this [J thinks the turtles are equidistant from the line of symmetry], but it's not going to go the right way.

Jodie has caught nicely two important mini-theorems: that symmetry would necessarily imply meeting on the mirror line; and the converse too - if they do not meet, then they cannot be symmetrical.

I would like to elaborate two points that emerge from this little example. The first centres on the ways that learners can use technology to shape their mathematical expression - how some elements of the invariant relationships between the given objects are identified and related within the symbolic discourse of the environment In the sense that these invariant relationships remain articulated only within activity using the notational system of the virtual world, they likely could not be said to constitute a formal abstraction But to the extent that they become transformed into something coherent, reusable and general, it does make sense to consider such activity as involving an abstraction of some kind (For further elaboration of this argument, in the context of stochastic thinking, see Pratt, 1998; Pratt and Noss, 2002 For a study in relation to students' conceptualisations of non-Euclidean geometry from a similar perspective, see Stevenson, 1996; Stevenson and Noss, 1999.)

The second point is related to the first, and concerns differential performance Put bluntly, children who may be apparently unable to express *any* relationships about their figures with pencil and paper are able to express them quite adequately (and sometimes quite elegantly) with the computer.

Reports of differential performance depending on context are commonplace. There are consistent and widely-reported findings concerning the differential performance between adults carrying out tasks in everyday settings and when given written assessments For example, Scribner's (1985) study of the dairy industry, Lave, Murtaugh and de la Rocha's (1984) investigation of weight-watchers, the seminal work of Saxe (1991) and Nunes, Schliemann and Carraher's (1993) investigations of street vendors have all shown convincingly that people who are error-prone in tests are mostly error-free in familiar practical contexts and that there is a major disjuncture between the strategies used in the two settings.

More generally, and especially since the work of Jean Lave and Etienne Wenger (1991) and others in broader anthropological contexts, we may more or less take for granted the situated view of knowledge genesis A key insight is that people construct solutions in the course of action and that these solutions are structured by activity In the supermarket, for example, Lave *et al.* (1984) illustrate how people avoid doing what might be classified as school mathematics not because it is too hard, but because the practice of supermarket shopping carries with it its own discourse and its own mechanisms for meaning-making One point that is often missed, however, is that we cannot conclude that there is nothing that passes for mathematical in shoppers' activities. The point is that when shoppers do use mathematics, it is supermarket mathematics, a mathematics made possible through the resources of the setting.

Since these studies, the situated cognition perspective has become ubiquitous. In its extreme version, it claims that "every cognitive act must be viewed as a specific response to a specific set of circumstances" (Resnick, 1991, p. 4) However, such arguments - compelling as they are - present researchers of mathematical learning with a number of seemingly intractable difficulties. If mathematics cannot be regarded as a decontextualised resource to be learned and then mapped onto settings, if it can only be defined in relation to specific situations, then we seem to have come close to distilling the mathematical essence out of mathematical thought.

One way out of this difficulty has been indicated by considering the role of representations. While seeking to understand the way schemas of 'if-then' reasoning were developed, Cheng and Holyoak (1985) posited the inadequacy of the specific-experience view and proposed the notion of *pragmatic reasoning schemas* which are guided as much by goals and purposes as the logic of the argument and are:

> primarily the products of induction from recurring experiences with classes of goal-related situations (p. 414)

The claim is that:

> people typically reason using abstract knowledge structures organised pragmatically, rather than in terms of purely syntactical rules of the sort that comprise standard logic. (Cheng *et al.*, 1986, p 314)

Similarly Nunes, Schliemann and Carraher (1993) show convincingly how street vendors could generalise their knowledge flexibly and argue that:

> the fact that specific information is contained in representations in street mathematics is not a drawback. It is specific information that allows subjects to control for the meaning and reasonableness of their answers in problem situations [...] Thus, representation of the particulars of a situation does not imply that the subject is restricted to understanding exact situations There is ample evidence for flexibility and generalisability of the pragmatic schemas of street mathematics. (p. 147)

These accounts concur on one important point: that individuals engaged in situated activity can and *do* generalise beyond the specificities of situations. Three questions naturally arise. What is it that is generalised? What is *abstract* about such knowledge? What is *situated* about such knowledge?

To try to answer these questions, Hoyles and I proposed, some ten or so years ago, the idea of *situated abstraction* (a first attempt is in Hoyles and Noss, 1992) as a tool to aid in understanding how learners construct mathematical ideas by drawing on the material and discursive components of a particular setting or, as we might put it today, activity system *Situated abstraction* seeks to describe metaphorically how a conceptualisation of mathematical knowledge can be both situated and abstract. It may be finely tuned to its constructive genesis - how it is learned, how it is discussed and communicated - and to its use in a cultural practice, yet can simultaneously retain mathematical invariants abstracted within that community of practice. The reflective symmetry example above gives a simple illustration of this idea

A traditional view regards an abstraction as containing some essential property of a situation extracted from it - not contained within it; it is deemed to be 'apart' from, even above, the situation of its genesis (see, for example, Piaget, 2000, p. 4). From this perspective, abstractions are, by definition, not embedded in situations. Rather, they involve expressions (internally or externally manifested) which, although they may derive from specific situations, are meant to represent a shift away from that situation. This perspective leads to the widely-shared assumption that practice is *merely* concrete and that abstraction is an extra-practical form of thought: see Wilensky (1991)

The idea of abstraction as a conceptualisation or a piece of knowledge lying in a separate realm from action, tools, language or indeed from any external referential sign system is important from a perspective of mathematical discourse, since mathematical discourse is normally conceived as self-contained: it forms part of a system that has its own objects and its own rules for transforming them.

This characteristic of formal mathematical abstraction is central to its utility: situated abstraction does not seek to challenge that utility, but questions whether mathematical abstractions can ever be fully separated from the context of their construction or application. Our broader hope is that the idea of situated abstraction will contribute to a theory of how mathematical knowledge is used or 'transferred' across settings (for other contributions to this emerging theory see, for example, Carraher and Schliemann, 2002; Sfard, 2001; Nemirovsky, in press; Hershkowitz, Schwarz and Dreyfus, 2001).

At the point at which Hoyles and I began to formulate these ideas, they were essentially hypotheses, based only on data derived from children and adults engaged with computationally expressive media. Fortunately, in the event, we were subsequently able to test these ideas in studies of mathematics in work and it is to these that I now turn.

Into the workplace

The study of mathematics in work provides a particularly fruitful setting for illuminating fundamental issues concerning the construction of mathematical meanings. Specifically, it affords an opportunity to focus on the situatedness of mathematical meanings by investigating their use rather than their genesis. In the long term, one might hope that such insight will afford leverage on the design of didactical strategies within and ultimately beyond work-based settings and I will outline some of these towards the end of this article.

Although it is clear that persons studied in their communities yield rich and useful data that describe what they do, it remains desirable to locate and elucidate the mathematical knowledge that they know. To achieve this aim, our group in London has employed ethnographic and interview data to capture meanings created *in situ* and the dialectical relationship of these meanings with mathematical expression on the one hand and professional expertise on the other. This has involved Celia Hoyles, Stefano Pozzi and myself in a series of studies with investment bank employees, paediatric nurses and commercial pilots; more recently, Phillip Kent and myself have been working with a group of structural engineers.

These professional groups differ in substantial ways, but there are similarities: in the explicitness of their mathematical training and in their intolerance - to a greater or lesser extent - of errors. We have developed a map of mathematical workplace activities comprised of documentary analysis, interviews with senior staff in each profession, general and task-based interviews with practitioner volunteers and ethnographic observation of these subjects in the workplace.

I will now try to summarise some of the outcomes of this research. I do so by sketching five vignettes, chosen to illustrate the outline of the theoretical position I wish to advance in the form of a set of provisional 'results'. The text of each vignette is based on the relevant co-authored papers that are referenced within it.

Vignette 1: The epistemological fragmentation of the workplace

The first vignette is drawn from a study conducted with a group of bank employees, part of which attempted to understand the bankers' ways of thinking about quantitative data better (see Noss and Hoyles, 1996a). In one of the task-based interviews, we presented the bankers with a problem of graphical interpretation, given in Figure 3 opposite.

The responses of the bank employees were surprisingly uniform. Most identified the graphs as a visual display of numbers, as a pictorial representation of underlying data rather than as a functional relationship and as an indication of a trend in data that allowed prediction: as one of them put it, a graph was the end result of a table of statistics that enabled him "to see [...] faster [...] the implications of the data". Where we saw graphs as a medium for expressing relationships (e.g. between quantity and time), bankers saw a display of data.

The origins of this epistemological diversity are almost certainly to be found in the tools of the system in which the bankers operate. On each employee's desk were several computers. Some, the traders and the operations staff, had three or four. On all but a very few screens, there were columns of data, graphs, and more columns of data: in every sense, graphs *were* pictures of numbers, rather than graphical representations of a functional relationship.

This epistemological standpoint with respect to graphical representations can, it seems, be thought of as the graphical face of a fragmented knowledge structure that characterises the practice of investment banking. We encountered departments specialising in the finest detail on one financial

An agent received commission for each transaction as follows:

(a) for transactions less than £30,000 ...	£750 plus $2\frac{1}{2}$% of the transaction.
(b) for transactions more than £30,000 ...	5% of the transaction

Which graph shows this situation most realistically?

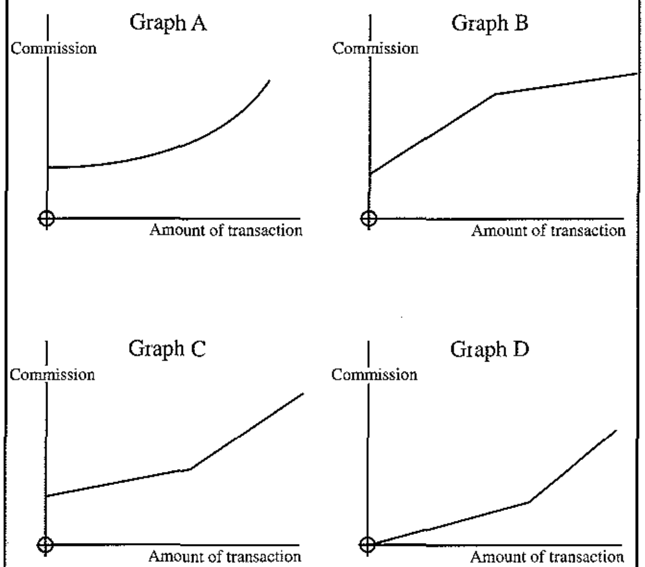

The correct graph is

Explain why.

Figure 3: Probing notions of quantity and time in relation to graphs: the correct option is Graph C

instrument, sharing a common wall but no common language with another - essentially similar - department. Of course, similarity is in the eyes of the beholder: while we might view, say, Nominal Certificates of Deposit and Treasury Bills as flavours of similar financial instruments sharing the same (or nearly the same) mathematical structure, the bankers saw finely-tuned pragmatic knowledge and strategies - and a discourse that served to reinforce the differences between them.

So the first major finding of our studies can be summarised as follows:

> *Result 1: There is an epistemological fragmentation of the knowledge structure of the workplace that shapes and is shaped by the discourse of the working practice. Strategies are finely tuned to the pragmatic demands of work activities, with little tendency to strive for a theoretical orientation involving generality or appreciation of unifying models*

Vignette 2: The role of artifacts and tools

The idea that people think and act within socio-cultural contexts which are mediated by cultural tools is now commonplace The work of Vygotsky, Luria and Leont'ev, indeed the entire corpus of work on activity theory, offers compelling evidence that both individual and social acts of problem solving are contingent upon structuring resources, involving a range of artifacts such as notational systems, physical and computational tools, and work protocols (Gagliardi, 1990).

Workplace settings are, naturally enough, littered with artifacts These artifacts are, for the most part, a simple expression of work protocols, so that in routine use – and the overwhelming majority of time in working practices is spent on routine – the structure of the artefact is hidden from view For example, in one study on a hospital ward (Pozzi, Noss and Hoyles, 1998), we found that a seemingly straightforward artefact like a fluid balance chart contained within it the crystallised activity (Leont'ev, 1978; see also Wertsch, 1985) of the hospital community, shaping in complex but mainly unnoticed ways the actions and discourse of those using it. A central part of this crystallised activity was a mathematical model of essential variables and relationships embedded in the activity: evidence for both the complexity and the invisibility of this mathematical model was gained by observing the ambiguity and uncertainty felt by a newcomer to the paediatric ward, as well as the extreme difficulty faced by the old-timers in communicating to her the structure that they had come to take for granted.

The arrival of the newcomer on the ward served to trigger a 'breakdown' or decision point within routine practice, a situation in which the models underpinning artifacts and the representational infrastructures on which their use depends rise to the surface and become open for inspection and negotiation by participants (and observation by researchers) That this model is normally hidden should cause no surprise: I have already noted that the purpose of an artefact is to facilitate the pragmatic activities of the workplace, not to learn mathematics or to gain insight into underlying models.

Nevertheless, when breakdowns do occur, invisible relationships buried in artifacts do not suffice and there is a need for the community to understand at least some of the workings of the models, to examine their strengths and limitations and to scrutinise the results of the mathematical labour congealed within them (see Hall, 1999, for a similar finding). At least in breakdown situations, we are abruptly made aware of circumstances that require more than mere procedural routine and the learning of work protocols. They require systemic interpretation – the individual is required to make sense of what she does within the confines of the broader socio-technical system

Result 2: Tools and artifacts shape activities and thought in ways that only become visible at times of breakdown to routine. In disruptions to routine, individuals need to develop a broader interpretative view of the model that underpins their routine practice.

Vignette 3: The anchoring of mathematical meanings in practice

It will help to focus on a specific knowledge domain I will turn to one of the most widely-researched topics in the field, ratio and proportion. Researchers on proportional reasoning in school and the workplace have distinguished two ubiquitous classes of strategies for making proportional calculations, *functional* (across measures) and *scalar* (within a single measure) – see Vergnaud (1983) for a thorough analysis. Nunes, Schliemann and Carraher (1993) have suggested that scalar strategies offer a mechanism for holding on to situational meaning by keeping only one measure in view. By way of contrast, functional strategies tend to be seen as semantically sparse manipulations of numerical quantities *per se*

It appears that this difference is the reason why people tend to prefer scalar strategies, even when they result in a more computationally awkward calculation. This is the crux of the counterposition in the literature of scalar and functional approaches, in that the privileging of the former has arisen from the apparent necessity in the latter to relinquish meaning in the form of a situational referent. Nunes, Schleimman and Carraher concluded that scalar approaches are drawn from experiences in everyday situations, are more flexible and generalisable than easily forgotten algorithmic approaches and, most relevant here, allow people to preserve the meaning of the situation by keeping variables separate and not calculating across measures.

To see how robust this finding is, I turn to the case of a group of paediatric nurses (see Hoyles, Noss and Pozzi, 2001) who are similarly expert in their field, but who have had years of school mathematical education as well as professional training. Nurses are drilled to perform accurate drug-dose calculations from an early stage in their professional training. One key aspect of this training is what, in our early interviews with them, several described as the 'nursing mantra' which states: "What you want, over what you've got, times the amount it comes in" or, in written form:

$$\frac{\text{What you want}}{\text{What you've got}} \times \text{The amount it comes in}$$

This rule is a version of the 'rule of three' and is completely general in scope Furthermore, it mirrors the actions that a nurse must take in order to prepare a prescription: look at the drug dose prescribed on the patient's chart ('what you want'); note the mass of the packaged drug to hand ('what you've got'), and then the volume of solution ('the amount it comes in'). This match of rule with actions and artifacts provides a possible explanation for the fact that at no time did we ever hear the nursing rule described in any other order.

During ethnographic observations and interviews, we noticed that while all the nurses' drug calculations were carried out correctly (unsurprisingly, in written tests, nurses' responses were highly error-prone), the strategies adopted were varied and exhibited a far richer complexity than would be suggested either from our interviews or from our expectations derived from the nursing literature. Of

30 episodes related to drug administration (out of a total of 250), we collected during 80 hours of observation, 26 combinations of ratios were observed with a variety of drugs, packaging and prescriptions Of these, only four involved the nursing mantra, while equal numbers chose scalar and functional strategies.

When we looked more closely at the strategies employed, we certainly found evidence of scalar strategies – as Vergnand and Nunes *et al.* would predict – even in the face of much simpler calculations being available with a functional approach. But we also found something surprising. I will illustrate with an example.

Belinda needed to give 120 mg of an antibiotic, amakacine, prepared in 100 mg per 2 ml vials. Before performing the calculation, she prepared for the administration and retrieved two vials of the antibiotic At this point, she found the volume she had to give with a fluency that was difficult to follow:

> *Belinda*: Amakacine [reads doses on the two vials] one hundred; one hundred; [reads year of expiry] ninety-eight; ninety-eight; [finds volume to be given] two point four mils.

A short interview with Belinda later revealed the nature of her strategy:

> *Int*: I didn't see you do any calculating there at all. You just drew it up. [...]
>
> *Belinda*: I knew the doses. [...] I know that that one is two point four . . . two point four mils. *With the amakacine, whatever the dose is, if you just double the dose, it's what the mil is* Don't ask me how it works, but it does. [emphasis added]
>
> *Int*: Why, what's the [...]?
>
> *Belinda*: One hundred and twenty mg, right [dose] and it comes in [...] and it goes in one hundred milligrams per two mils. So if you double it, that makes two hundred and forty [...] two point four mils
>
> *Int*: I'm sorry I don't understand
>
> *Belinda*: So if you just double it up. Double one twenty; one twenty and one twenty is two hundred and forty. And the dose is two forty. So very often *that's how it is with amakacine,* so if you're giving eighty [...] eighty milligrams to give, and if you double it up, it's one point six. [emphasis added]

Belinda's description clearly indicates a transformation from the dose mass to the dose volume, so in this sense the strategy is functional. But a simple classification of the strategy as functional does not do it justice. Her description suggests that the operation was associated with the drug itself rather than with the ratio between the mass and volume: "That's how it is with amakacine", says Belinda, apparently seeing the allowable arithmetic operation and the particular drug itself as intimately connected. Similarly, her description of the strategy suggests that she was neither simply manipulating numbers (or even quantities) nor performing arithmetic operations on them Rather, she described the transformation as 'doubling up' and effortlessly combined into a single process what would generally be recognized as the doubling operation and the movement of the decimal point

In this episode, we see an illustration of how the nurses often opted for strategies that would, in the literature, be described as unlikely and lacking in meaning. Our interpretation of these findings is that the nurses' knowledge of concentration, that is their appreciation of the invariance of the relationship between mass and volume as evidenced in their drug calculations, was anchored in an intimate knowledge of the drug itself, as well as in the properties of familiar packaging constraints of prescribed doses The knowledge was mutually constituted and expressed as both mathematical relation and culturally-shared situational noise – the same kind of knowledge that we encountered earlier in the context of computer worlds and which we called situated abstraction.

> *Result 3: Knowledge is mutually constituted by a co-ordination produced in activity of mathematical knowledge and situational noise to form situated abstractions.*

Vignette 4: The qualitative restructuring of mathematical knowledge in activity

In a recent study, Phillip Kent and myself have been investigating the ways in which mathematical knowledge is conceived and deployed with employees of a large London-based engineering firm (see Kent and Noss, 2001, 2002) We have encountered, even with this mathematically educated group, a ubiquitous view that the majority of structural engineers do not 'use mathematics' of any sophistication in their professional careers. So, while all believed that it was important for graduate engineers to have an appreciation for advanced mathematics, it is something they would rarely be expected to use.

> Once you've left university you don't use the maths you learnt there, 'squared' or 'cubed' is the most complex thing you do. For the vast majority of the engineers in this firm, an awful lot of the mathematics they were taught, I won't say learnt, doesn't surface again.

I think that this particular engineer's description of mathematics as not 'surfacing' is a fortuitous one. We have seen in the case of the nurses that mathematical knowledge becomes fused with professional knowledge as situated abstraction, not as abstraction in its pure form. But it is this pure form, particularly for mathematically sophisticated groups such as engineers, that is readily recognisable as mathematics. Our engineer is right that mathematics does not surface; or rather, that it seldom surfaces in the form it was learned and taught. It has been transformed into something else, something at once more usable, more embedded, more noisy. Only the vestigial traces of the college mathematics taught to engineers remains in the mathematics that they actually use in activity. [2]

The transformation in the character of mathematics appears to be not simply a quantitative one, nor merely a replacement of mathematical activity by professional expertise and experience. It represents a qualitative, epistemological and cognitive restructuring of the mathematics as it becomes 'embedded' in engineering expertise [3] I claim that engineers' conceptualisations of this restructured mathematical knowledge are legitimately considered as situated abstractions.

I will illustrate with an example. The type of qualitative thinking that characterises the use of 'feel' in the engineering design process is exemplified by the concept of *load path*, the notion that the loads acting on a structure have to 'flow down into the ground' like a kind of fluid. It is a powerful, very physical concept and extremely useful because it provides a way of thinking about a structure before any analysis is done, allowing judgements to be made about the validity of a quantitative analysis of the structure. One engineer put it thus:

> A load is applied and eventually it's got to get back into the ground. It's so fundamental to structural design that you have to be able to see what that route is in order to have a feeling, to be able to calculate, what sorts of loads and forces will be apparent in any particular member Without a clear idea of the load path, you have nothing to judge what you're getting from the computer

Formal mathematical analysis, on the other hand, is based on the assumption of static equilibrium, which assumes that nothing is moving in a stable structure, an assumption that appears to conflict with the load-path concept Nevertheless, for the engineer above, load path has become a situated abstraction of stability criteria: it allows predictions of behaviour that emerge from fusing together the actual properties of the material (e g. steel beams) with the associated (mathematically-abstracted) forces (see also Bissell and Dillon, 2000)

The attribution of 'mythical' chains of causality to formally non-causal situations has been studied by researchers in various areas of cognitive science, although not, as yet, within the context of mathematics education Viewed in this light, it is tempting to view this as an idiosyncrasy, a technique that works well enough in practice, as an approximation or useful metaphorical approximation to the 'real' mathematical abstraction ($\sum Forces = 0$). In fact, it turns out that the view of force as momentum flow has a long and epistemologically coherent pedigree. diSessa (1980) makes a compelling case for a view of force as:

> simply the flow of the conserved 'stuff', momentum, from one place to another. (p 2)

He notes that broadly convergent views have been advanced by no lesser scientific figures than Mach, Kirchoff and Hertz [4]

For the moment, the relevant point is this: engineering discourse employs, in at least one important way, a kind of knowledge which is at once about mathematical relations and about substance The idea of flow makes no sense without something to flow through - the beams and struts of everyday engineering practice. Mathematical knowledge has been transformed (in this case, there is an epistemological isomorphism) to the extent that even those engaged in it do not necessarily recognise its existence. This poses sharply two questions: how does the formally-learned knowledge (e g the engineers' knowledge of Newton's laws, or the nurses' knowledge of the drug-dose mantra) become transformed, both cognitively and culturally, into something new and more functional within professional practice and what connection, if any, is maintained between them?

I have no data on these questions For the moment, the key issue concerns the transformation of knowledge, the creation of new epistemologies as a characteristic part of professional expertise. Here, at least, is the explanation of the apparent invisibility of mathematical activity. Here, too, is a broader, more culturally oriented perspective on the hitherto individualistic notion of situated abstraction that recognises the individual's embedding in an ambient social and cultural space

> *Result 4: As mathematical knowledge is embedded in new settings and activities, it undergoes an epistemological and cognitive transformation What is consciously thought of as mathematics by practitioners appears to be only the visible component of a larger, transformed body of mathematics in use that takes the form of situated abstractions*

Vignette 5: The situatedness of abstraction

The final vignette will deal with the most problematic (and so far, under-researched) issue. The challenge is to test the situatedness of knowledge, to assess the extent to which knowledge in the form of situated abstraction 'transfers' to new situations (or better still, to find a convincing alternative metaphor for the notion of transfer itself)

For this task, I will return to the nurses (this vignette is extracted from Noss, Hoyles and Pozzi, 2002). Across many different drug administrations, with varying degrees of complexity and in a variety of situations, nurses' procedures for calculating drug dosages consistently retained a constant covariation of mass and volume in the drug solution: this is the epistemological core of the situated abstraction of concentration. In order to probe the nature of this knowledge, it was necessary to devise a methodology that could tease out the limits and situatedness of abstractions developed in activity. Our solution was to displace the nurses from their familiar practice, by 'forcing' them to reflect and articulate what it was they knew, and how - if at all - they thought about it in relation to their practice We did this by a series of task-based interviews, in which nurses were progressively faced with situations that were further and further removed from the practices we had observed, yet retained elements of familiar situations for them.

I will summarise the findings. First, when the nurses were faced with a close simulation of their practice, they displayed similar strategies to those identified in the ethnographic studies, together with a strong sense of the invariant relationship of mass and volume. In these cases, the nurses' reasoning was supported by a synergy of their existing (school) mathematical knowledge and their practical experience. Thus, knowledge of the invariance of drug concentration characterised nurses' strategies even when they were removed from their practice; they engaged with the underlying objects of mass, volume, concentration and rate and the relationships among them, in order to develop effective strategies, such as mentally dividing extensive quantities into visualisable chunks in ways that made a direct connection to the artifacts of their practice.

Second, and by contrast, an analysis of the nurses' responses to a less familiar scenario illustrated that when it became impossible to link contextual elements with mathematical knowledge, the nurses' responses became far less clear. Our conclusion was that it was crucial for the nurses to exploit the texture of their experience as a resource in their mathematical activity: when the texture of that experience became unavailable, the mutually constitutive elements of professional and mathematical knowledge became disconnected.

It would be difficult to explain the nurses' situated yet abstract knowledge, if it were merely to consist of a collection of abstract procedures or, conversely, if it were entirely contingent on participation in the specificities of nursing practice – that is, if the mobilisation of the nurses' knowledge depended wholly upon immersion in the cultural practice that gave rise to it. Moreover, it is clear that the noise of the situation is a critical element of the conceptualisation of the mathematical knowledge used in the practice, one that affords the extension of a situated abstraction into less familiar and novel domains. It also limits its generalisability.

> *Result 5 (conjecture): The noise of a situation forms a core part of a situated abstraction. When it can be called upon in a new situation (and only then?), the mathematical knowledge can be 'transferred'*

Designing for change

I promised at the outset to draw out implications of these work studies and to elicit some general principles concerning the design of mathematical practices for learning. The hypothesis is that the ways in which people reconstruct knowledge for use in work is spontaneous, in the sense of it deriving from participation in the practices of the community and, for the most part, not being formally taught within the practice. That being so, I might further hypothesise that, given the functionality of this kind of knowledge, one might attempt to design and construct activity systems for learning that harness the features of the workplace, at least those that are perceived as constitutive of learning. So I will try, very briefly, to map the set of findings from the maths-in-work studies onto a set of implications for design. This will also provide an opportunity to refocus this article back on the roles of digital technologies and to begin to deliver on the undertaking I gave at the outset in this regard.

The first finding (vignette 1) concerned the fragmentation of knowledge within the ecology of the workplace system. If it is more generally true that strategies are pragmatically oriented, perhaps we should design environments which explicitly and visibly demonstrate the power of (mathematical) invariants. This power is singularly lacking as an explicit focus of most mathematics curricula. And, as we have seen, the role and function of invariants is hardly a natural priority in the world of work. This observation adds something to the 'real world' movement that pervades – at least at a rhetorical level – the stated aims and methodologies of various curricula. It focuses attention on the construction of models of reality (rather than reality itself), an activity in which the identification of what does and what does not vary, as well as how, is a crucial component. As I will clarify below, this is an initial point of contact with the particularly powerful role that digital technologies can play.

The second finding (vignette 2) was that knowledge is pervasively structured by artifacts. Artifacts-in-activity – or more properly, the knowledge congealed within them – do much of the work involved in understanding and predicting the behaviour of the workplace, as part of a distributed system of knowledge construction within it. Yet learning environments are typically spartan in their use of artifacts. I conjecture that lots of manipulable, combinable and *useful* things are a key part of realising the design challenges we face. Here, too, we will see the special contribution that digital technologies can make: the addition of a digital dimension to learning-oriented artifacts can be exploited to increase the range of expressive power and creative possibilities afforded by the manipulation of artifacts and, crucially, the potential to isolate and reflect upon the mechanisms and models that endow them with functionality.

Two further findings (vignettes 3 and 5) concerned the role of situated abstractions. We saw that situated abstractions are mutually constituted by mathematical knowledge and situational noise, and that situated abstractions extend to new situations to the extent that the contextual 'noise' of their genesis can be carried alongside formal, mathematical knowledge. Two possible implications arise. First, that we should include noise as a carefully-designed element of learning environments, not as contextual mess (a particularly irritating practice in the U.K. is to wrap any mathematical idea in contextual clutter and label it 'situated'), but as culturally-shared situations that are meaningful for its participants. Second, we need to design systems that afford the construction of new situations from old ones, in ways that allow the knowledge-constitutive elements of noise to remain invariant. This approach would, I think, represent rather a radical shift in much current pedagogical practice: it suggests that the standard modelling strategy of removing noise in order to expose underlying structures might fruitfully be rethought in favour of a view of contextual noise as an element of what makes knowledge learnable, functional and, for want of a better word, transferable.

A final lesson (vignette 4) from the mathematics-in-work studies involved the transformation of mathematical knowledge as it crosses boundaries of an individual's experience, or its cultural embedding in different types of work situations. This is perhaps the most difficult finding from which to draw a canonical implication for design, but I will choose just one I conjecture that the shifting character of mathematical meanings within different representational infrastructures is endemic to the workplace and, more generally, beyond it If that is the case, we would do well to consider designing systems that afford a range of representational systems for expressing mathematical relationships, rather than focusing simply on one, standard infrastructure that has evolved over time for purposes other than pedagogical utility

I have summarised the findings and their implications schematically in Figure 4. This summary should be taken as merely indicative, as any map from the workplace findings onto a set of design principles can hardly claim uniqueness, perhaps achieving at best an outline research agenda rather than a list of definitive implications

	If .	then we should design to .
1	knowledge is fragmented and strategies pragmatic,	demonstrate the power of invariants.
2a	knowledge is pervasively structured by artifacts and their underlying models,	supply lots of Really Useful Things
2b	people need to understand the models,	make things that people can see inside.
3	situated abstractions are mutually constituted by mathematical knowledge and situational noise,	maximise situational noise in culturally-relevant ways.
4	mathematical knowledge is transformed when it structures new activities,	respect the mathematical epistemologies of new representational forms.
5	situated abstractions depend on noise for 'transfer',	afford construction of new situations from old ones.

Figure 4: Some schematic implications of the mathematics in work findings for the design of learnable environments

In drawing some tentative implications from the maths-in-work studies, I have mentioned the computational presence several times. It is now time to concentrate explicitly on the place of digital technologies. In our book *Windows on Mathematical Meanings: Learning, Cultures and Computers* (Noss and Hoyles, 1996b), Celia Hoyles and I argue that constructing runnable models in the form of computer programs affords a compelling example of a learnable mathematics, opening unique opportunities for students to interact with a formal system. In modifying or constructing a model of a system, a student must articulate rigorously its salient relationships, describing mathematical structures in a language that can be communicated, extended and become the subject of reflection.

There are many advocates of a similar perspective (see Hoyles and Noss, in press a, for a review). In a recent study, for example, Sherin (2001) proposes that programming-based representations might be easier for students to understand physics with than equation-based representations and that programming-based representations might privilege a somewhat different 'intuitive vocabulary', i.e. might tap into different things that people 'just know'. I would add a third point: that programming can afford access to a rich and extensible set of situated abstractions of physical relationships that I think correspond to what he calls *a physics of processes and causation* (as opposed to algebra-physics which he characterises as *a physics of balance and equilibrium*)

It is not important whether we accept Sherin's strong conjecture or not: in *Windows*, we refer to LogoMathematics or Programming Mathematics to emphasise that it *is* a different kind of mathematics that is at issue (this is an instance of the fourth design challenge). What is important is that we recognise that the switch from one representational form to another carries with it the possibility of a simultaneous switch in both epistemology and learnability.

I would like to underline a further element of the consideration of design principles for learning environments: the importance of mathematical models, a proximally social issue that tacitly underpins much of what has gone before I have not dealt with this problem in any depth, except in noting (in vignette 2) that in breakdown situations, individuals at work are explicitly required to interpret and understand elements of the models that underpin the artifacts and work systems they otherwise take for granted.

I believe the knowledge economy has massively broadened the number of people who need to understand the system they are using: elsewhere (Noss, 1998, 2002), I elaborate a case that competence in constructing, interpreting and critiquing models has become a core part of social and professional life in the twenty-first century. As profit margins are squeezed, and globalisation intensifies, the fall-out of the knowledge economy applies to greater slices of the (first-world) workforce – not to everyone, but to substantial and increasing sections of it.

Not many individuals need constantly to access the precise details of the models that underpin their social and professional existence, but I am convinced there are more than is evident at first sight. Models are genuinely pervasive; more and more people need to know what a model *is*, even if they cannot build one; to understand what a variable *is*, even if they cannot write the relevant equation that defines it; to interpret the output (and inputs) of a model, even if they cannot grasp the model as a unified whole.

Vast sections of the workforce operate with models every day – in the form of spreadsheets – even if most of their workings are purposefully obscured in 'macros' or in opaquely encoded recipes governing their use.

Sharing, critiquing and representing models is massively under-represented in mathematics curricula, still wedded to the epistemological and pedagogical requisites of the nineteenth century rather than transforming both in the face of the demands and computational possibilities of the twenty-first Elsewhere (Noss, 2002), I have labelled the needs of the knowledge economy as requiring a *meta-epistemological stance*:

- knowing *that* things work in programmed ways rather than (necessarily) *how*;
- knowing *that* there are assumptions instantiated in the choice of variables and that there are relationships between them;
- knowing *about* connections between variables rather than calculational knowledge about their detailed interrelationships;
- knowing *about* interpreting and critiquing models, together with the different representational forms in which they may be expressed.

This stance is also about the ways in which this kind of knowledge is communicated to others who interact with other parts of the same system or other, linked systems (see also Kaput, Hoyles and Noss, 2002)

In short, I contend that manipulating, modifying, constructing and sharing computationally instantiated models of mathematical systems affords the best chance we have for designing a more learnable mathematics and of realising the five challenges outlined in the previous section.

My colleagues and I have recently completed a study aimed at instantiating this approach in the *Playground Project* [5] *Playground* has involved a group of researchers based in four European countries who have developed a system with which young children, aged less than eight years old, can play, share, construct and rebuild computer games. Our goal has been to put children in the role of game designers and game programmers, rather than merely consumers of games programmed by adults, and to engage them in exploring and understanding the formal rule systems that underpin game play and game design. Our broader, hitherto untested, belief is that the children's deep engagement with a formal system of this kind will serve as a powerful generic knowledge substrate on which future mathematical learning might be based

I will do no more than sketch an outline of the design of our *Playground*. [6] The learner is confronted with a world in which things happen and, more importantly, can be made to happen. It is full of objects - balls, spaceships, characters, balloons and many, many more. Most of these have properties, or behaviours - they bounce, fly, walk, talk, and so on. They interact: when the spaceship is hit by a balloon, it may explode; when a ball hits an edge, it can bounce and makes a bang. Finding out how an object 'works' is straightforward: one simply flips it over and inspects the program, amending it to one's taste.

Programs are not lines of text or even icons; programs are animated robots, who have been trained by being given an example to remember. The act of programming consists of giving a robot a set of objects and a sequence of actions to perform on it: she remembers both (in her thought bubble) and applies the sequence to any set of objects that matches those with which she was trained. Abstraction (how to generalise from a given instance) is achieved not by introducing variables, but by erasing specificity

There is much more to the design of *Playground* than the above paragraph can possibly convey (see http://www.ioe.ac.uk/playground for a comprehensive overview; see also, Hoyles and Noss, in press b). I will resist elaborating the design of the project and its findings here; equally, I will leave as an exercise to the interested reader the various ways in which the design of the project 'conforms' to the design principles I outlined above. The latter enterprise, while elegant, would presuppose a much more detailed elaboration of the environment and the learning outcomes associated with it. Instead, I will focus on a single issue that, among the many raised by the *Playground* study (as well as computational environments in general), returns us to the issue of mathematical epistemology that has formed an underlying theme for this article

Consider the case of a child designing a game fragment in which a ball is to be made to move across the screen as the mouse is moved (the case sketched here is based on a real episode with an eight-year-old boy, reported in Noss, 2002). How should that movement be instantiated? One way is to borrow the behaviour of some other, pre-existing object that already has a similar behaviour: perhaps there is a nearby spaceship whose behaviour can be copied and pasted (in *Playground*, these 'system' actions are performed by animated characters, not by key presses). Pasting the spaceship's behaviour onto the ball has the desired result or at least near enough for a first attempt. But it is not quite right; some fine tuning is necessary and this, in turn, provokes some engagement with the program that makes the spaceship (and now the ball) work. It turns out (let us say) that the two-dimensional motion of the ball is instantiated as the vector sum of horizontal and vertical components. Of course, it does not look that way to the child: it might, for example, be that there are two robots (one called *'move left and right'* and the other called *'move up and down'*).

Think for a moment of the knowledge congealed in the innocent phrase 'vector sum'. Concealed in this phrase is a taken-for-granted representational infrastructure that includes the definition of a vector, the algebraic system for combining two or more vectors and a range of properties (e.g. scalar and vector product) that give meaning to the very idea of what a vector is and why it is a conceptually powerful generalisation of a real number. This structure is relatively complex and is postponed with good reason until the latter stages of compulsory education, if it is taught there at all. Yet the complexity is in the infrastructure, not the idea.

The point is that what is and what is not intuitive is hugely contingent on the representational infrastructure with which the intuition is expressed. In the *Playground*, the addition of vectors is instantiated not as an algebraic relation but as a natural property of the representational system. The (object-oriented) structures of the system translate, more or less directly, into what kinds of things can be taken for granted as being 'just so', what meanings can be derived from them, and most importantly, the ways in which the objects and their programmed behaviours can be made functional within a given situation. In short, the representational infrastructure has transformed not only the learnability of the mathematical knowledge, but the mathematical epistemology *at work* in the activity system.

Concluding remarks

This last point brings me to the intention I flagged at the outset, to conclude with the notion of epistemology *at* rather than *in* work. What is the connection between the two? A key link is that the analysis of mathematics in work concerns the transformation of knowledge as it is recontextualized across settings. We have seen how a person's mathematical knowledge is not invariant across time and space; it is transformed into different guises, different epistemologies, more or less visible in the form of mathematics, as the map of an individual's participation in new activity systems is continually redrawn. This transformation seems much more powerful than the traditional notions of 'application' or 'use' that is often employed as a metaphor to describe this process. I have argued that recognising these transformations and designing learning environments that exploit them is a priority for the construction of a more learnable mathematical epistemology.

More generally, I have elaborated a further point of connection between cognitive and cultural perspectives. In imagining how mathematical structures can be externalised and manipulated within an appropriately expressive representional structure, I have indicated how abstractions constructed within concrete situations may compensate for their lack of universality by their gain in expressiveness. When general relationships can be expressed, they can be explored and become familiar. In the process, the links with knowledge of lived-in cultures can be maintained, rather than severed in the quest for ultimate pinnacles of abstraction

The objects that populate *Playground* are every bit as concrete and real to a learner as the load path on the components of a bridge are to an engineer. Like their professional counterparts, children are engaged in an activity that researchers in the field of mathematical learning may recognise as having a mathematical component, but which are to the child merely part of the ecological system - the totality of relationships between themselves and the environment and the ways in which these are expressed and communicated.

That the *Playground* and mathematical epistemologies run side-by-side should not be a matter of surprise: there is, after all, no single way in which humans can conceptualise their environment (mathematically or otherwise), even though some are socially and historically privileged within a given culture. Official, symbolic mathematics is privileged in just this way and there are good reasons for this. But the compactness and elegance of mathematical expression does not necessarily make it equally functional for learning and, if learning is our prior goal, we would do well to think about new epistemological frameworks in which to embed the mathematics we wish our students to understand. New epistemologies mean new intuitions, new things to be built with them and new means for combining and reconstructing them. They involve new sets of situated abstractions that are both functional and powerful I think this is the major challenge for the design of didactical environments, to create new systems which might, I think, be justifiably described as involving new mathematical epistemologies at work.

Acknowledgements

The corpus of data on which this paper is built was the result of several research projects and some ten or more years of collaboration with colleagues I acknowledge the support of the Economic and Social Research Council, U.K. (grants R 0222 50004 and R 0002 23420) and the European Union (grant 29329). Almost all of the data presented here has been published in collaboratively-written articles, which are appropriately referenced My thanks to Ross Adamson, Miki Grahame, Lulu Healy, Celia Hoyles, Phillip Kent, Dave Pratt and Stefano Pozzi for their collaboration in the various projects. My thanks to Lulu and Celia who offered helpful comments on an earlier draft. *Windows on Mathematical Meanings* has provided a source for the main theoretical ideas, and in this respect - as in so many others - I would like to acknowledge the profound help and support of my co-author and collaborator, Celia Hoyles.

Notes

[1] This paper is an expanded and revised version of a plenary address to the Twenty-sixth International Conference on the Psychology of Mathematics Education, University of East Anglia, June 2002

[2] In fact there is a division of labour in engineering into 'analysts' and 'designers'. The former – who represent a few percent of the profession – do indeed use mathematics in an explicit and readily recognisable way

[3] Michèle Artigue (in press) has made a related point deriving from the work of Yves Chevallard

[4] diSessa (1980) suggests that a view of force as momentum flow may more easily engage and refine students' existing intuitions and therefore present a more learnable physics than that represented by the familiar $F = ma$

[5] The *Playground* project invovled a consortium across four countries, directed by myself and Celia Hoyles. The London team also comprised (at various times) Ross Adamson, Miki Grahame, Sarah Lowe and Dave Pratt Ken Kahn, the author of *ToonTalk*, was a consultant to the project

[6] Actually, there are two 'playgrounds'; the second, *Pathways*, will not be referred to here. See, for example, Goldstein, Noss, Kalas and Pratt (2001)

References

Artigue, M (in press) 'Learning mathematics in a CAS environment: the genesis of a reflection about instrumentation and the dialectics between technical and conceptual work', *International Journal of Computers for Mathematical Learning.*

Bissell, C. and Dillon, C. (2000) 'Telling tales: models, stories and meanings', *For the Learning of Mathematics* **20**(3), 3-11

Carraher, D and Schliemann, A. (2002) 'The transfer dilemma', *Journal of the Learning Sciences* **11**(1), 1-24.

Cheng, P., Holyoak, K., Nisbett, R. and Oliver, L (1986) 'Pragmatic versus syntactic approaches to training deductive reasoning', *Cognitive Pyschology* **18**(3), 293-328.

Cheng, P. and Holyoak, K. (1985) 'Pragmatic reasoning schemas', *Cognitive Psychology* **17**(4), 391-416.

Cobb, P and Bowers, J (1999) 'Cognitive and situated learning perspectives in theory and practice', *Educational Researcher* **28**(2), 4-15.

diSessa, A. (1980) *Momentum Flow as an Alternative Perspective in Elementary Mechanics*, Division for Study and Research in Education, Cambridge, MA, MIT

diSessa, A (1987) 'Artificial worlds and real experience', in Lawler, R. and Yazdani, M (eds), *Artificial Intelligence and Education. Volume One: Learning Environments and Tutoring Systems*, Norwood. NJ, Ablex Publishing Corporation, pp. 55-78

Gagliardi, P. (1990) *Symbols and Artifacts: Views of the Corporate Landscape,* New York, NY, Aldine de Gruyter

Goldstein, R., Noss, R., Kalas, I and Pratt, D (2001) 'Building rules', in Beynon, M, Nehaniv, C and Dautenhahn, K (eds), *Proceedings of the 4th International Conference of Cognitive Technology CT2001,* Coventry, Warwickshire, University of Warwick, pp. 267-281

Hall, R (1999) 'Following mathematical practices in design-oriented work', in Hoyles, C., Morgan, C. and Woodhouse, G. (eds), *Rethinking the Mathematics Curriculum*, London, Falmer Press, pp 29-47

Healy, L. (2002) *Iterative Design and Comparison of Learning Systems for Reflection in Two Dimensions,* Unpublished Ph.D thesis, London, Institute of Education, University of London

Hershkowitz, R., Schwarz, B. and Dreyfus, T (2001) 'Abstraction in context: epistemic actions, *Journal for Research in Mathematics Education* **32**(2), 195-222.

Hoyles, C. and Noss, R. (1992) 'Looking back and looking forward', in Hoyles C and Noss R (eds), *Learning Mathematics and Logo,* Cambridge, MA, MIT Press, pp. 431-468

Hoyles, C., Noss R and Pozzi, S (2001) 'Proportional reasoning in nursing practice', *Journal for Research in Mathematics Education* **32**(1), 4-27

Hoyles, C. and Noss, R. (in press a) 'What can digital technologies take from and bring to research in mathematics education?', in Bishop, A (ed.), *Second International Handbook of Mathematics Education* Dordrecht, Kluwer Academic Press.

Hoyles, C and Noss, R. (in press b) 'Playing with rules: the collaborative construction of an adventure game', in Blatchford, J (ed), *Experimental School Environments,* Cambridge, Cambridge University Press

Kaput, J, Noss, R. and Hoyles, C (2002) 'Developing new notations for a learnable mathematics in the computational era', in English, L (ed.), *Handbook of International Research in Mathematics Education* Hillsdale, NJ, Lawrence Erlbaum, pp 51-75.

Kent, P. and Noss, R (2001) 'The mathematical components of engineering expertise: the relationship between doing and understanding mathematics', *Proceedings of the IEE Second Annual Symposium on Engineering Education: Professional Engineering Scenarios,* **2**, pp 39/1 - 39/7.

Kent, P. and Noss, R. (2002) *The Mathematical Components of Engineering Expertise: Final Report to the ESRC* (http://www ioe ac.uk/moss/MCEE)

Kieran, C., Forman, E and Sfard, A. (2001) 'Guest editorial, *Educational Studies in Mathematics* **46** (1-3), 1-11

Lave, J., and Wenger, E (1991) *Situated Learning: Legitimate Peripheral Participation,* Cambridge, Cambridge University Press

Lave, J., Murtaugh, M. and de la Rocha, O. (1984) 'The dialectic of arithmetic in grocery shopping', in Rogoff, B. and Lave, J (eds), *Everyday Cognition: Its Development in Social Context* Cambridge, MA, Harvard University Press, pp 67-94.

Leont'ev, A. (1978) *Activity, Consciousness and Personality* (trans Hall, M.), Englewood Cliffs, NJ, Prentice-Hall

Nemirovsky, R. (in press) 'On guessing the essential thing', in Gravemeijer, K, Lehrer, R., van Oers, B and Verschaffel L. (eds), *Symbolizing, Modeling and Tool Use in Mathematics,* Dordrecht, Kluwer Academic Publishers

Noss, R (1998) 'New numeracies for a technological culture' *For the Learning of Mathematics* **18**(2), 2-12.

Noss, R (2002) 'For a learnable mathematics in the digital culture, *Educational Studies in Mathematics*

Noss R. and Hoyles, C. (1996a) 'The visibility of meanings: modelling the mathematics of banking', *International Journal of Computers for Mathematical Learning* **1**(1), 3-31.

Noss, R. and Hoyles, C (1996b) *Windows on Mathematical Meanings: Learning, Cultures and Computers,* Dordrecht, Kluwer

Noss, R., Hoyles, C and Pozzi, S. (2002) 'Abstraction in expertise: a study of nurses' conceptions of concentration', *Journal for Research in Mathematics Education,* **33**(3), 204-229

Nunes, T., Schliemann, A. and Carraher, D. (1993) *Street Mathematics and School Mathematics* Cambridge, Cambridge University Press

Piaget, J (2000) *Studies of Reflecting Abstraction* (ed and trans Campbell, R) London, Taylor and Francis.

Pozzi S., Noss R and Hoyles C. (1998) 'Tools in practice, mathematics in use', *Educational Studies in Mathematics* **36**(2), 105-122

Pratt, D. (1998) *The Construction of Meanings in and for a Stochastic Domain of Abstraction,* Unpublished Ph.D Thesis, London, University of London.

Pratt, D., and Noss, R. (2002) 'The micro-evolution of mathematical knowledge: the case of randomness, *Journal of the Learning Sciences,* **11**(4), 453-488.

Resnick, L. (1991) 'Shared cognition: thinking as a social practice', in Resnick, L., Levine, J. and Teasley, S. (eds), *Perspectives on Socially Shared Cognition,* Washington, DC, American Psychological Association, pp. 1-20

Saxe, G (1991), *Culture and Cognitive Development: Studies in Mathematical Understanding,* Hillsdale, NJ, Lawrence Erlbaum.

Scribner, S (1984) 'Studying working intelligence', in Rogoff, B.and Lave, J. (eds), *Everyday Cognition. Its Development in Social Context,* Cambridge, MA, Harvard University Press, pp. 9-40.

Sfard A (2001) 'There is more to discourse than meets the ears: looking at thinking as communicating to learn more about mathematical learning', *Educational Studies in Mathematics* **46**(1-3), 13-57

Sherin, B (2001) 'A comparison of programming and algebraic notation as expressive languages for physics', *International Journal of Computers for Mathematical Learning* **6**(1), 1-61

Smith, J., diSessa, A. and Roschelle, J. (1993) 'Misconceptions reconceived: a constructivist analysis of knowledge in transition', *Journal of the Learning Sciences* **3**(2), 115-163

Stevenson, I (1996) *The Iterative Design of a Computer-based Microworld for Non-Euclidean Geometry* Unpublished PhD thesis, London, University of London.

Stevenson, I and Noss, R. (1999) 'Supporting the evolution of mathematical meanings: the case of non-Euclidean geometry', *International Journal of Computers for Mathematical Learning* **3**(3), 229-254.

Vergnaud, G (1983) 'Multiplicative structures', in Lesh, R and Landau, M (eds), *Acquisition of Mathematics Concepts and Processes* London, Academic Press, pp. 127-174.

Wertsch, J. (ed) (1985) *Culture, Communication and Cognition: Vygotskian Perspectives,* New York, NY, Cambridge University Press.

Wilensky, U. (1991) 'Abstract mediations on the concrete and concrete implications for mathematics education', in Harel, I and Papert, S (eds), *Constructionism,* Norwood, NJ, Ablex Publishing Corporation, pp 193-203